

COSMIC COVARIANCE AND THE LOW QUADRUPOLE ANISOTROPY OF THE WILKINSON MICROWAVE ANISOTROPY PROBE (WMAP) DATA

LUNG-YIH CHIANG¹, PAVEL D. NASELSKY^{2,3}, PETER COLES⁴

Submitted to the Astrophysical Journal

ABSTRACT

The quadrupole power of cosmic microwave background (CMB) temperature anisotropies seen in the WMAP data is puzzlingly low. In this paper we demonstrate that Minimum Variance Optimization (MVO), a technique used by many authors (including the WMAP science team) to separate the CMB from contaminating foregrounds, has the effect of forcing the extracted CMB map to have zero statistical correlation with the foreground emission. Over an ensemble of universes the true CMB and foreground are indeed expected to be uncorrelated, but any particular sky pattern (such as the one we happen to observe) will generate non-zero measured correlations simply by chance. We call this effect “cosmic covariance” and it is a possible source of bias in the CMB maps cleaned using the MVO technique. We show that the presence of cosmic covariance is expected to artificially suppress the variance of the Internal Linear Combination (ILC) map obtained via MVO. It also propagates into the multipole expansion of the ILC map, generating a quadrupole deficit with more than 90% confidence. Since we do not know the CMB and the foregrounds *a priori*, there is therefore an unknown contribution to the uncertainty in the measured quadrupole power, over and above the usual cosmic variance contribution. Using the MVO on a series of Monte Carlo simulations that assume Gaussian CMB fluctuations, we estimate that the real quadrupole power of the CMB lies in the range $[305.16, 400.40](\mu\text{K}^2)$ (at the $1 - \sigma$ level).

Subject headings: cosmology: cosmic microwave background — cosmology: observations — methods: data analysis

1. INTRODUCTION

Detailed measurements of the pattern of temperature anisotropies in the cosmic microwave background (CMB) have provided cosmologists with unprecedented opportunities to probe the physics of the early Universe. In particular, the Wilkinson Microwave Anisotropy Probe (WMAP) data (Bennett et al. 2003b,c; Hinshaw et al. 2003b, 2007) have played a pivotal role in the establishment of the standard “concordance” cosmological model and have heavily constrained alternatives to the inflationary paradigm for the origin of large-scale structure.

According to the concordance cosmology, the variations in temperature of the CMB across the celestial sphere should possess a very broad spectrum usually expressed in terms of the amplitude of spherical harmonic modes labelled by the usual angular harmonic frequency ℓ . The power of the pattern of anisotropies power is expected to be strongest for the quadrupole ($\ell = 2$) which correspond to variations on an angular scale of 90° on the sky. In principle, therefore, the quadrupole should be the easiest harmonic mode to detect. In practice, however, the measured quadrupole anisotropy of the CMB from WMAP seems to sit rather uncomfortably with the standard cosmology: its amplitude seems too low (Efstathiou

2003a; de Oliveira-Costa et al. 2004). Of course, the power that resides in each multipole of the CMB pattern (including the quadrupole) is not so much measured as extracted. There are various diffuse foreground emissions to be separated from the “true” primordial CMB, such as free-free emission arising from electron-ion scattering, synchrotron emission from cosmic ray electrons accelerating in the galactic magnetic fields, and dust emission. However, although a number of authors have employed different foreground subtraction schemes, all find the derived quadrupole power to be lower than the Concordance Cosmological Model with a significance level of 5% or better (Efstathiou 2003a; Eriksen et al. 2004; de Oliveira-Costa & Tegmark 2006; Park, Park & Gott 2007; Saha et al. 2007). This has even prompted suggestions of physics beyond the standard model (Efstathiou 2003b; Luminet et al. 2003).

In this paper we suggest that the quadrupole power deficit in WMAP may be an artifact of the Minimum Variance Optimization method used by many authors to clean CMB maps from foreground contamination. In the next Section, we explain the basics of the MVO method. In Section 3 we show how this method is prone to a bias introduced by the accidental alignment of structures in the CMB and foreground templates leading to an effect we call “cosmic covariance”. In Section 4 we quantify the likely effect of this bias using Monte Carlo simulations, and we estimate the quadrupole power taking into account the effect of cosmic covariance. We present our conclusions in Section 6.

2. MINIMUM VARIANCE OPTIMIZATION

Although various schemes have been proposed for subtracting foregrounds from CMB data, one concept that

Electronic address: lychiang@asiaa.sinica.edu.tw, naselsky@nbi.dk, peter.coles@astro.cf.ac.uk

¹ Institute of Astronomy and Astrophysics, Academia Sinica, P.O. Box 23-141, Taipei 10617, Taiwan, R.O.C.

² Niels Bohr Institute, 17 Blegdamsvej, Copenhagen, DK-2100, Denmark

³ Space Research Department, Southern Federal University, Zorge, 5, 344091, Rostov on Don, Russia

⁴ School of Physics & Astronomy, Cardiff University, 5 The Parade, Cardiff, CF24 3AA, Wales, United Kingdom

tends to be in common for multi-frequency cleaning methods is Minimum Variance Optimization (MVO). The theoretical basis for MVO is the primordial CMB temperature anisotropies arise from black-body radiation and therefore constitute a frequency-independent signal that persists across different observed frequency bands (Tegmark & Efstathiou 1996). The observed sky temperature in a given band i is therefore just

$$T_i = T_c + f_i, \quad (1)$$

where T_c is the CMB signal and f_i the foreground emission band i , which will be a composite of dust, free-free and synchrotron sources. For the purposes of this paper we assume that instrument noise is negligible. The frequency-independent signal can be flushed out by an Internal Linear Combination (ILC) from the maps of M frequency bands with weighting coefficients w_i , where $\sum_{i=1}^M w_i = 1$ (Bennett et al. 2003c):

$$T_{\text{ilc}}(p) = \sum_{i=1}^M w_i T_i(p) = T_c(p) + \sum_{i=1}^M w_i f_i(p) \equiv T_c(p) + T_{\text{fr}}(p), \quad (2)$$

where $T_{\text{ilc}}(p)$ denotes the ILC map value at a pixel p and T_{fr} the foreground residual. In general,

$$\text{Var}(T_{\text{ilc}}) = \text{Var}(T_c) + \text{Var}(T_{\text{fr}}) + 2\text{Cov}(T_c, T_{\text{fr}}), \quad (3)$$

where the variances and covariance are measured over an ensemble of possible skies. The MVO technique is based on the *a priori* assumption that the frequency-independent signal (i.e. the true CMB) should be *statistically* independent of any foreground contribution, so that the covariance term vanishes and

$$\text{Var}(T_{\text{ilc}}) = \text{Var}(T_c) + \text{Var}(T_{\text{fr}}), \quad (4)$$

or in other words,

$$\langle \sigma_{\text{ilc}}^2 \rangle = \langle \sigma_c^2 \rangle + \text{Var} \left[\sum_{i=1}^M w_i f_i(p) \right], \quad (5)$$

where the angle brackets denote ensemble averages as above. Minimizing the variance of the ILC map is then equivalent to minimizing the foreground contamination of a map of thermal noise.

The WMAP 3-year ILC map and, more importantly, the power spectrum for $\ell \leq 30$ are⁵ thus produced by employing the MVO on 5 frequency maps K (23 GHz), Ka (33 GHz), Q (41 GHz), V (61 GHz) and W (94 GHz) bands in the pixel domain for 12 separate regions, where Region 0 marks the largest region with $|b| \geq 15^\circ$ (about 89% of the whole sky), and Regions 1 – 11 are those around the Galactic plane. For each region the 5 frequency band maps are linearly combined and a set of weighting coefficients $w_i^{(R)}$ for each region are obtained in such a way that the resultant variance is forced to be minimum:

$$\frac{\partial \sigma_{\text{ilc}}^2}{\partial w_i} = \frac{\partial}{\partial w_i} \text{Var} \left[\sum_{i=1}^M w_i f_i(p) \right] = 0. \quad (6)$$

Note that Equation (5) involves the requirement of statistical independence over an ensemble of universes. As

⁵ For power spectrum for $\ell > 30$, WMAP science team applies a χ^2 foreground template fitting method.

we have only one universe (i.e. one CMB sky), the calculation of variances and covariances can only be done over the set of pixels corresponding to a single sky, and not over an entire probability distribution. The sky covariance between T_c and T_{fr} need not be exactly zero, since there is always a chance of some chance coincidence of features such as hotspots between foreground and true CMB. If this happens then

$$\frac{\partial \sigma_{\text{ilc}}^2}{\partial w_i} \neq \frac{\partial}{\partial w_i} \text{Var} \left[\sum_{i=1}^M w_i f_i(p) \right]. \quad (7)$$

As we shall now show, this leads to a bias in the recovery of the CMB map.

3. COSMIC COVARIANCE

Let us assume (for illustrative purposes only) that the components of the foreground emission at different frequencies have the same spatial distribution (i.e uniform foreground spectra). In reality the foregrounds at different frequency bands have spatial variations of the spectral indices, which can amplify the resulting bias. In the uniform case we have $f_i(p) = S_i F(p)$, where $S_i \equiv S(\nu_i)$ is the composite frequency spectrum of the foreground emission and F is the common spatial distribution over the set of pixels. Thus frequency band maps $T_i(p) = T_c(p) + S_i F(p)$ and the ILC map

$$T_{\text{ilc}}(p) = T_c(p) + \Gamma F(p), \quad (8)$$

where $\Gamma \equiv \sum_i w_i S_i$. The variance of the ILC map can be written as follows (Hinshaw et al. 2007)

$$\sigma_{\text{ilc}}^2 = \langle T_{\text{ilc}}^2(p) \rangle - \langle T_{\text{ilc}}(p) \rangle^2 = \sigma_c^2 + 2\Gamma \sigma_{c,F} + \Gamma^2 \sigma_F^2, \quad (9)$$

where

$$\sigma_{c,F} \equiv \langle T_c F \rangle - \langle T_c \rangle \langle F \rangle \quad (10)$$

is the “cosmic covariance” between the CMB and the foreground spatial distribution F . In this and the subsequent equations, the angle brackets denote averages over the pixels of a single sky rather than ensemble averages; c.f. Equation (5).

The term “cosmic covariance” is coined similarly to “cosmic variance” (White, Krauss and Silk 1993) because both effects emanate from the fact that we must do our statistical analysis using a single version of the celestial sphere. Cosmic variance arises from the fact that a single sky does not offer sufficient spherical harmonic modes to make an accurate estimate of the ensemble-averaged power at low multipoles. In the case of cosmic covariance, the single sky prevents us from being sure that our calculation of the correlation between foreground and CMB is exact, particularly for the *quadratic minimization*. The value $\sigma_{c,F}$ will be non-zero merely because the CMB and the foregrounds happen to line up to a certain degree simply by chance. If both CMB and foreground contain large-scale structures then the effect of cosmic covariance will be larger as the sky will contain fewer independent regions.

The cosmic covariance $\sigma_{c,F}$ can also be expressed via the correlation coefficient $X_{c,F}$: $\sigma_{c,F} \equiv X_{c,F} \sigma_c \sigma_F$, where $X_{c,F}$, the “cosmic correlation coefficient” characterizes the level of resemblance in morphology between the CMB and the foreground spatial distribution, and $-1 \leq X_{c,F} \leq 1$.

	K	Ka	Q	V	W
K		0.9986	0.9968	0.9964	0.9778
Ka	0.9987		0.9994	0.9989	0.9749
Q	0.9966	0.9994		0.9991	0.9720
V	0.9957	0.9985	0.9993		0.9774
W	0.9835	0.9790	0.9756	0.9790	

TABLE 1

CORRELATION COEFFICIENTS X BETWEEN THE FULL-SKY WMAP FOREGROUND MAPS (LOWER LEFT TRIANGLE) AND BETWEEN THOSE IN THE WMAP-DEFINED REGION 0, WHICH ACCOUNTS FOR 89% OF THE FULL SKY (UPPER RIGHT TRIANGLE). ONE CAN SEE THAT K, Ka, Q AND V CHANNELS (EXCEPT W) HAVE CORRELATION VERY CLOSE TO UNITY. ALL OTHER REGIONS OF THE WMAP CHANNELS EXCEPT W SHOW THE SAME TREND.

Employing the MVO criterion to the variance of the ILC map, we get

$$\frac{\partial \sigma_{\text{ilc}}^2}{\partial w_i} = 2 \frac{\partial \Gamma}{\partial w_i} X_{\text{cF}} \sigma_{\text{c}} \sigma_{\text{F}} + 2 \Gamma \frac{\partial \Gamma}{\partial w_i} \sigma_{\text{F}}^2 = 0, \quad (11)$$

which gives $\Gamma = -X_{\text{cF}} \sigma_{\text{c}} / \sigma_{\text{F}}$ and the ILC map is then

$$T_{\text{ilc}}(p) = T_{\text{c}}(p) - \frac{X_{\text{cF}} \sigma_{\text{c}}}{\sigma_{\text{F}}} F(p). \quad (12)$$

One can see if cosmic covariance is zero (so is X_{cF}), then $T_{\text{ilc}}(p) = T_{\text{c}}(p)$.

The last term $-X_{\text{cF}} F(p) \sigma_{\text{c}} / \sigma_{\text{F}}$ is usually discussed in terms of an anti-correlation bias (Hinshaw et al. 2007), but it is in fact due to the MVO serving to eliminate the covariance between the ILC and the foreground spatial distribution. To see this, note that the covariance between the ILC and the foreground is forced to vanish:

$$\sigma_{\text{ilc}, \text{F}} \equiv \langle T_{\text{ilc}} F \rangle - \langle T_{\text{ilc}} \rangle \langle F \rangle = \sigma_{\text{c}, \text{F}} - \frac{X_{\text{cF}} \sigma_{\text{c}}}{\sigma_{\text{F}}} \sigma_{\text{F}}^2 = 0. \quad (13)$$

The ILC map is thus guaranteed to have a morphology that has no resemblance to that of the foregrounds, regardless of whether or not the CMB has any such resemblance. That is to say, even if there exists any cosmic covariance between the true CMB signal and the foregrounds, it is duly subtracted by the MVO, and the resulting ILC map must display zero sky correlation with the foregrounds.

Forcible subtraction of the unknown cosmic covariance also affects the overall variance. Putting $\Gamma = -X_{\text{cF}} \sigma_{\text{c}} / \sigma_{\text{F}}$ into Eq.(9), we have

$$\sigma_{\text{ilc}}^2 = \sigma_{\text{c}}^2 (1 - X_{\text{cF}}^2). \quad (14)$$

Subtraction of the cosmic covariance manifests itself as a deficit in the total variance of a factor $(1 - X_{\text{cF}}^2)$. This can partially account for the low variance in the CMB map (Monteserin et al. 2007). As we do not know T_{c} and F *a priori*, we are left with no information about the level of cosmic covariance after it is subtracted.

One of the dangers of the MVO approach is that it encourages circular reasoning. The *a priori* assumption is that the CMB and the foregrounds are statistically independent. The MVO produces an ILC map that has zero sky correlation with the foregrounds. This seems to confirm the correctness of this assumption. However, the MVO must produce this result whatever the morphology of the input templates.

Although the above analysis is based on the existence of composite foregrounds with uniform frequency spectra, we can see that it is not far from reality. In Table 1, we show the pixel-by-pixel correlation coefficients




FIG. 1.— Decreasing of correlation with the foreground due to the MVO. We show the histogram of X_{cF} and X_{iF} , the correlation coefficient between the foreground and the simulated CMB (left) and the retrieved CMB (right) via the MVO.

between the WMAP composite foregrounds in WMAP-defined Region 0, which contains $\sim 89\%$ of the whole sky (upper right triangle). One can see that all channels have excellent agreement in terms of morphology (with the exception of W channel due to synchrotron emission). The lower left triangle part shows correlations between full-sky foreground maps, which also display the same trend.

4. MONTE CARLO SIMULATIONS

To demonstrate the power deficit and reduction of cosmic covariance in the ILC map due to the MVO, we perform simulations of 2000 sets of frequency band maps corresponding to the WMAP K, Ka, Q, V and W channels. These comprise the sum of a simulated CMB signal (assuming Gaussian fluctuations), mock foreground maps and instrument noise maps. The CMB signal is simulated with the WMAP best-fit Λ CDM model. The mock foreground maps at WMAP 5 frequency bands are produced as follows: we take the WMAP composite foreground at K band, which is multiplied by a factor ρ_i such that $\rho_i \sigma_{\text{K}} = \sigma_i$, $i = \text{Ka, Q, V and W}$ where σ_i is the standard deviation of the WMAP composite foreground maps (in thermodynamic temperature units). As such, these mock foreground maps have exactly the same morphology, but each has the same variance as the WMAP foreground map of the corresponding frequency band. The Gaussian instrument noise map is with $C_{\ell} = 2 \times 10^{-8}$ (mK²). To follow the WMAP procedure, both the simulated CMB and noise maps are smoothed to 1° FWHM before summation with the mock foregrounds (which are already smoothed). The MVO is then employed on the internal linear combination of the band maps to retrieve the CMB signal.

In Figure 1 we plot the histogram of X_{cF} (left), the correlation coefficient between the simulated CMB and the common foreground map, and that of X_{iF} (right), the correlation coefficient between the retrieved CMB (ILC) and the common foreground map. One can see the sky correlation is reduced below a level 2×10^{-4} after the application of the MVO.

The deficit in the variance $X_{\text{cF}}^2 \sigma_{\text{c}}^2$ is registered in the angular power spectrum C_{ℓ} via $\sigma_{\text{ilc}}^2 = (4\pi)^{-1} \sum_{\ell} (2\ell + 1) C_{\ell}^{\text{ilc}}$. Note that C_{ℓ} is always positive, so the deficit in variance manifests itself as a few dips compared to the desired CMB angular power spectrum C_{ℓ}^{c} . Since the MVO acts on the overall variance, it only ensures the variance is minimum; the result for individual spherical harmonic modes is less obvious. Decomposing Eq.(12) into spherical harmonic coefficients $a_{\ell m}^{\text{ilc}} = a_{\ell m}^{\text{c}} - X_{\text{cF}} a_{\ell m}^{\text{F}} / \sigma_{\text{F}}$, we




FIG. 2.— Normalized histograms (thin) and cumulative histograms (thick) of the power difference $\Delta(\delta T_\ell^2) \equiv \delta T_{\ell,i}^2 - \delta T_{\ell,s}^2$ from 2000 realizations for $\ell = 2, 3, 4$ and 5 , where $\delta T_{\ell,s}^2$ denotes the simulated CMB power at ℓ mode and $\delta T_{\ell,i}^2$ the retrieved CMB (ILC) power. The MVO is employed on frequency band maps with the foregrounds having the same morphology (see text for the simulations). One can see that the CMB retrieved via MVO has a skew distribution for $\ell = 2$, indicating the power of the retrieved CMB is smaller than the simulated ones.

$$C_\ell^{\text{ilc}} = \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} |a_{\ell m}^{\text{ilc}}|^2 = C_\ell^c + \frac{X^2 \sigma_c^2}{\sigma_F^2} C_\ell^F - \frac{2X\sigma_c}{\sigma_F(2\ell+1)} \sum_{m=-\ell}^{\ell} |a_{\ell m}^c| |a_{\ell m}^F| \cos(\phi_{\ell m}^c - \phi_{\ell m}^F) \quad (15)$$

One can see that it is possible for some multipoles to acquire excess power (bumps) over C_ℓ^c if, for instance, $\pi/2 < \phi_{\ell m}^c - \phi_{\ell m}^F < 3\pi/2$. So in order to preserve the overall deficit, it shall drag the dips at other multipoles even further down.

In Figure 2 we plot the normalized histogram and cumulative histogram of the power difference $\Delta(\delta T_\ell^2) \equiv \delta T_{\ell,i}^2 - \delta T_{\ell,s}^2$, where $\delta T_\ell^2 = \ell(\ell+1)C_\ell/(2\pi)$ is the temperature anisotropy power at multipole ℓ , the subscripts s and i denote the simulated CMB and the retrieved CMB from MVO, respectively. It is easy to see that, particularly for multipole number $\ell = 2$, the simulation ensemble forms a skewed distribution with a tail on the negative side, indicating the power of the retrieved CMB is more likely to be smaller than that of the input (simulated) CMB. For other multipoles the skew is less obvious. This shows, therefore, that the deficit in variance due to the MVO is most likely registered in the quadrupole.

One should note that the power deficit claimed by the WMAP team is calculated from Region 0 (Full sky with Galaxy cut) and then corrected by using a maximum likelihood estimate (Hinshaw et al. 2007). The power we have calculated above, however, is directly from full-sky simulations because one can then single out the effect of the MVO, whereas extrapolation of the power from incomplete sky coverage may introduce yet another systematic error, particularly for large angular scales.

5. ESTIMATION OF THE QUADRUPOLE POWER

Since we know the cosmic covariance is going to be subtracted by the MVO, there is an unknown error in the quadrupole power even before the cosmic variance becomes involved in the interpretation of the data. Consequently, the quadrupole power we can estimate is at




FIG. 3.— Distribution of the CMB quadrupole power. We employ the MVO on frequency band maps of the WMAP K, Ka, Q, V channels (based on the excellent agreement in morphology between the foregrounds) and add on the estimated estimated quadrupole the missing power from 10^4 realizations of 4-channel MVO. The true CMB quadrupole falls in $[305.16, 400.40](\mu K^2)$ (at the $1 - \sigma$ level).

best not going to be a single value, but a sampling distribution which we construct with help of Monte Carlo simulations, assuming the CMB is a Gaussian random field. According to Table 1 (lower left triangle), the full-sky WMAP foregrounds at K, Ka, Q and V channels (except W) have correlation coefficients deviated from unity by less than 5×10^{-3} . We can employ the MVO on these 4 channels to retrieve the quadrupole power and then add on the missing power from Monte Carlo simulations.

The weighting coefficients for the linear combination from the 4 frequency band maps are $(6.2419 \times 10^{-3}, -0.21597, -0.37970, 1.5894)$ and we denote the resultant map as ILC4. The correlation coefficients between the ILC4 and the derived foregrounds (frequency band maps subtracting ILC4) at K, Ka, Q and V are $X = (8.3553, 8.3510, 8.3351, 8.2826) \times 10^{-8}$, respectively, which confirms the presence of negligible correlations after the application of MVO. The quadrupole power from the ILC4 map is $\delta T_2^2 = 326.30 \mu K^2$. We then add the missing power for quadrupole from 10^4 realizations of exactly the same procedure as described earlier in this Section, except this time using WMAP Q band foreground as common foreground for the 4 channels. In Fig.3 we plot the histogram of the resulting estimates of true CMB quadrupole power, which falls in $[305.16, 400.40](\mu K^2)$ (at the $1 - \sigma$ level) .

6. DISCUSSION AND CONCLUSIONS

We have shown that the MVO method for eliminating foreground contamination from CMB maps is prone to a bias if there exist accidental alignments between foreground features and structures in the true CMB sky. This bias acts in such a way that a suppression of the quadrupole is more likely than an enhancement. It also reduces the overall variance of the recovered map compared to the true CMB sky. This suggests that attempts to attribute the low quadrupole to exotic physics may be premature.

The MVO was originally designed for the separation of Gaussian signals, whose statistical properties are completely characterized by their power spectra. For a linear combination of such signals, the variance is clearly the natural choice for a parameter to be minimized, and in such a case this method generally leads to a good recon-

struction of the power in each component. In this paper, we have demonstrated even for the MVO to work on Gaussian signals it requires a statistical ensemble: for a single realization the role of cosmic covariance becomes crucial. The foregrounds relevant to CMB analysis are also highly non-Gaussian, which makes this problem even worse.

For composite foregrounds with varying spectra, the variance after application of the MVO will induce even more foreground contamination in the ILC map because the weighting coefficients will then be tuned not only for the subtraction of cosmic covariance, but for varying morphology among the different foregrounds. It is

then possible that the variance deficit generated by the subtraction of cosmic covariance can be compensated by a foreground residual, which can push the quadrupole power back to the value that is consistent with the Concordance Cosmological model. Extra caution, therefore, will have to be applied when employing the MVO to the interpretation of results from the upcoming ESA Planck Experiment.

Acknowledgments: We acknowledge the use of HEALPIX⁶ package (Górski, Hivon & Wandelt 1999) to produce $a_{\ell m}$ from the WMAP data, and the use of GLESP⁷ package.

⁶ <http://healpix.jpl.nasa.gov/>

⁷ <http://www.glesp.nbi.dk/>

REFERENCES

- Bennett, C. L., et al., 2003, ApJS, 148, 1
 Bennett, C. L., et al., 2003, ApJS, 148, 97
 de Oliveira-Costa, A., Tegmark, M., 2006, Phys. Rev. D, 74, 023005
 de Oliveira-Costa, A., Tegmark, M., Zaldarriaga, M., Hamilton, A., 2004, Phys. Rev. D, 69, 063516
 Efstathiou, G., 2003, MNRAS, 343, L95
 Efstathiou, G., 2003, MNRAS, 346, L26
 Efstathiou, G., 2004, MNRAS, 348, 885
 Eriksen, H K., Banday, A J., Górski, K M., Lilje, P B., 2004, ApJ, 612, 633
 Górski, K. M., Hivon, E., Wandelt, B. D., 1999, in A. J. Banday, R. S. Sheth and L. Da Costa, Proceedings of the MPA/ESO Cosmology Conference “Evolution of Large-Scale Structure”, PrintPartners Ipskamp, NL
 Hinshaw, G., et al., 2003, ApJS, 148, 135
 Hinshaw, G., et al., 2007, ApJS, 170, 288
 Luminet, J.-P., Weeks, J., Riazuelo, A., Lehoucq, R., Uzan, J.-P., 2003, Nature, 425, 593
 Monteserin, C., et al., 2007, MNRAS submitted (arXiv:0706.4289)
 Park, C.-G., Park, Changbom., Gott, J R III., 2007, ApJ, 660, 959
 Saha, R., Prunet, S., Jain, P., Souradeep, T., 2007, Phys. Rev. D submitted (arXiv:0706.3567)
 Tegmark, M., 1998, ApJ, 502, 1
 Tegmark, M., de Oliveira-Costa, A., Hamilton, A., 2004, Phys. Rev. D, 68, 123523
 Tegmark, M., Efstathiou, G., 1996, MNRAS, 281, 1297
 White, M., Krauss, L M., Silk, J., 1993, ApJ, 418, 535